arXiv:1603.04572v1 [stat.ML] 15 Mar 2016

On the exact recovery of sparse signals via conic relaxations

Hongbo Dong*

March 16, 2016

Abstract

In this note we compare two recently proposed semidefinite relaxations for the sparse linear regression problem by Pilanci, Wainwright and El Ghaoui ("*Sparse learning via boolean relaxations*", 2015) and Dong, Chen and Linderoth ("*Relaxation vs. Regularization: A conic optimization perspective of statistical variable selection*", 2015). We focus on the cardinality constrained formulation, and prove that the relaxation proposed by Dong, etc. is theoretically no weaker than the one proposed by Pilanci, etc. Therefore any sufficient conditions of exact recovery derived by Pilanci can be readily applied to the Dong's relaxation, including their results on high probability recovery for Gaussian ensemble. Finally we provide empirical evidence that Dong's relaxation requires much fewer observations to guarantee the recovery of true support.

1 Two convex relaxations for sparse linear regression

Given a collection of observed sample points $(x_i, y_i) \in \mathbb{R}^p \times \mathscr{Y}$, the goal of a sparse learning task is to learn a linear function $x \mapsto \beta^T x$ that is then used to predict an outcome of $y \in \mathscr{Y}$ for future/unseen data, where β is restricted to have a small number of nonzero entries). Such a task can be modeled as the following cardinality constrained optimization problem

$$\min_{\substack{\beta\in\mathbb{R}^p,\\ \|\beta\|_0\le k}} \frac{1}{n}\sum_{i=1}^{n} f\left(\beta^T x_i; y_i\right). \tag{1}$$

With the cardinality constraint, (1) is usually highly nonconvex and difficulty to solve to global optimality. The authors in [2] considered the following regularized version,

$$\min_{\substack{\beta\in\mathbb{R}^p,\\ \|\beta\|_0\le k}} \frac{1}{n}\sum_{i=1}^{n} f\left(\beta^T x_i; y_i\right) + \rho\|\beta\|_2^2. \tag{2}$$

*Department of Mathematics and Statistics, Washington State University, Pullman, WA 99163

One of the key results in [2] shows that (2) can be equivalently formulated as minimizing a convex function over a subset of binary vectors,

$$\min_{\substack{z\in\{0,1\}^p,\\ \sum_j z_j\le k}} \underbrace{\max_{v\in\mathbb{R}^n}\left\{-\frac{1}{2\rho}v^T X\mathbf{D}(z)X^T v-\sum_{i=1}^n f^*(v_i;y_i)\right\}}_{G(z)}, \tag{3}$$

where $G(z)$ is convex because as it is the max function of infinite many linear functions, and $f^*(v,y):=\sup_{t\in\mathbb{R}}\{st-f(t,y)\}$ is the conjugate function of $f(\cdot;y)$.

In this note we focus on the important special case of sparse linear regression, i.e., we consider the following cardinality-constrained quadratic program,

$$v_{\ell 0}:=\min_{\beta}\ \frac{1}{2}\|X\beta-y\|_2^2+\frac{1}{2}\rho\|\beta\|_2^2,\ \ s.t.\ \|\beta\|_0\le k, \tag{$\ell_0:card$}$$

The authors of [2] further proposed to relax the binary condition $z\in\{0,1\}^p$ in (3) to $z\in[0,1]^p$, and studied the conditions under which such a relaxation is exact. When specialized to the sparse linear regression problem, the continuous relaxation takes the following form of a semidefinite program,

$$\begin{aligned} v_{PWG}:=\min_{t\in\Re,z\in[0,1]^p}\ & 0.5t \\ s.t.\ & \begin{bmatrix} t & y \\ y & I_n+\frac{1}{\rho}X\mathbf{D}(z)X^T\end{bmatrix}\succeq 0,\quad e^T z\le k, \end{aligned} \tag{SDP_{PWG}}$$

where e is a vector with all entries 1 in proper dimension, and $\mathbf{D}(z)$ is a diagonal matrix whose entries are $z_i, i=1,...,p$. It can also be equivalently written as the following compact form,

$$v_{PWG}=\frac{1}{2}\min_{z\in[0,1]^p,e^Tz\le k}\ y^T\left(\frac{1}{\rho}X\mathbf{D}(z)X^T+I_n\right)^{-1}y.$$

Following a different approach, authors of [1] recently proposed another semidefinite relaxation for sparse linear regression where the ℓ-0 norm appears as a regularized term. When modified as a convex relaxation for the cardinality constrained form (3), their proposed semidefinite relaxation is,

$$\begin{aligned} v_{DCL}:=\min_{b\in\mathbb{R}^p,B\in\mathscr{S}^p}\ & \frac{1}{2}\left\langle\begin{bmatrix} y^Ty & -y^TX \\ -X^Ty & \rho I_p+X^TX\end{bmatrix},\begin{bmatrix}1 & b^T\\ b & B\end{bmatrix}\right\rangle \\ s.t.\ & \begin{bmatrix}1 & b^T\\ b & B\end{bmatrix}\succeq 0 \\ & \begin{bmatrix} z_i & b_i\\ b_i & B_{ii}\end{bmatrix}\succeq 0,\forall i,\ \sum_{i=1}^p z_i\le k. \end{aligned} \tag{SDP_{DCL}}$$

In this note we compare these two semidefinite relaxations. We show that the relaxation (SDP_{DCL}) is no weaker than (SDP_{PWG}) in this section. In section 2 we establish a result that characterizes a certificate of exactness for the convex relaxation (SDP_{DCL}), hence extends a key result in [2] to (SDP_{DCL}). Section 3 concerns the probability of exact recovery for the case of Gaussian ensemble, where we show empirically (SDP_{DCL}) can recover the true support of with much less data points.

We first state a technical lemma that will be used soon.
Lemma 1. *For any $X \in \Re^{n\times p}$ and $\rho > 0$, we have*

$$\min_{\beta\in\mathbb{R}^p}\left\{\frac{1}{2}\|X\beta - y\|_2^2 + \frac{1}{2}\rho\|\beta\|_2^2\right\} \;=\; \frac{1}{2}y^T\left(\frac{1}{\rho}XX^T + I_n\right)^{-1} y$$

Proof. Straightforward computation. □

By Lemma 1, (SDP_{PWG}) can be reformulated as

$$\nu_{PWG} = \min_{z\in[0,1]^p, e^T z\le k}\ \min_{\beta\in\mathbb{R}^p}\ \frac{1}{2}\left\|X\sqrt{\mathbf{D}(z)}\beta - y\right\|_2^2 + \frac{1}{2}\rho\left\|\beta\right\|_2^2, \tag{4}$$

where $\sqrt{\mathbf{D}(z)}$ is a diagonal matrix with the i-th diagonal entry $\sqrt{z_i}$.
Proposition 1. $\nu_{\ell 0} \ge \nu_{DCL} \ge \nu_{PWG}$.

Proof. Suppose that $(\tilde{b}, \tilde{B}, \tilde{z})$ is optimal in (SDP_{PWG}). Without loss of generality we may assume that $\tilde{z}_i = \frac{\tilde{b}_i^2}{\tilde{B}_{ii}}$ for all $\tilde{B}_{ii} \neq 0$, and $\tilde{z}_i = 0$ otherwise. Therefore $\tilde{z}_i \in [0,1], \forall\, i$. Define $\tilde{\beta}$ as

$$\tilde{\beta}_i = \begin{cases} \tilde{z}_i^{-\frac{1}{2}}\tilde{b}_i, & if\ \tilde{z}_i > 0 \\ 0 & if\ \tilde{z}_i = 0. \end{cases}$$

Then obviously $\tilde{b} = \sqrt{\mathbf{D}(\tilde{z})}\tilde{\beta}$ and $(\tilde{\beta}, \tilde{z})$ is feasible in (4). We have

$$\begin{aligned}\frac{1}{2}\left\langle\begin{bmatrix} y^Ty & -y^TX \\ -X^Ty & \rho I_n + XX^T\end{bmatrix}, \begin{bmatrix}1 & \tilde{b}^T \\ \tilde{b} & \tilde{B}\end{bmatrix}\right\rangle &\ge \frac{1}{2}\left\langle\begin{bmatrix} y^Ty & -y^TX \\ -X^Ty & XX^T\end{bmatrix}, \begin{bmatrix}1 & \tilde{b}^T \\ \tilde{b} & \tilde{b}\tilde{b}^T\end{bmatrix}\right\rangle + \frac{1}{2}\rho\,\mathbf{trace}(\tilde{B}) \\ &\ge \frac{1}{2}\left\|X\sqrt{\mathbf{D}(\tilde{z})}\tilde{\beta} - y\right\|_2^2 + \frac{1}{2}\rho\left\|\tilde{\beta}\right\|_2^2 \ge \nu_{PWG}.\end{aligned}$$

The first inequality is because of $\tilde{B} \succeq \tilde{b}\tilde{b}^T$. The second inequality is because of $\tilde{B}_{ii}\tilde{z}_i \ge \tilde{b}_i$, which implies $\tilde{B}_{ii} \ge \tilde{\beta}_i^2$, and the final inequality is by the characterization (4). □

As (SDP_{DCL}) satisfies the Slater condition, strong duality holds and the dual of (SDP_{DCL}) is

$$
\begin{aligned}
v_{DCL} = \frac{1}{2} y^T y + \max_{\tau,\lambda,t,d} \; & -\frac{1}{2}\tau - \frac{1}{2} k\lambda \\
s.t. \; & \begin{bmatrix} \tau & -y^T X - t^T \\ -X^T y - t & X^T X + \rho I - \mathbf{D}(d) \end{bmatrix} \succeq 0 \qquad (SDP_{DCL} : dual) \\
& \begin{bmatrix} \lambda & t_i \\ t_i & d_i \end{bmatrix} \succeq 0, \forall i.
\end{aligned}
$$

2 Certificate of exactness

Proposition 1 implies that if $v_{PWG} = v_{\ell 0}$, then $v_{DCL} = v_{\ell 0}$. Therefore all sufficient conditions for the exactness of (SDP_{PWG}) readily carry over to (SDP_{DCL}). Authors of [2] provided a characterization of a *certificate of exactness* for the continuous relaxation of (3), as well as a specialized result on (SDP_{PWG}). We restate their characterization result in Theorem 1, and provide a parallel result in Theorem 2 for (SDP_{DCL}).

Theorem 1 (Corollary 2 in [2])**.** *The convex relaxation (SDP_{PWG}) is exact if and only if there is a subset $S \subseteq \{1, ..., p\}$, where $|S| \leq k$, such that there exists $\lambda \in \mathbb{R}_+$,*

$$
|X_j^T M y| > \lambda, \;\; \forall j \in S, \text{ and} \tag{5}
$$

$$
|X_j^T M y| \leq \lambda, \;\; \forall j \notin S, \tag{6}
$$

where $X_j \in \mathbb{R}^n$ is the j-th column of X, and $M := \left(I_n + \rho^{-1} X_S X_S^T\right)^{-1}$.

Using this result, the authors were able to prove a high-probability exact recovery condition for the special case of Gaussian ensembles. We leave the discussion of Gaussian ensemble in the next section. Here we provide a parallel characterization of certificates of exactness for (SDP_{DCL}).

Theorem 2. *Let $S \subseteq \{1, ..., n\}$, $|S| = k$ and z^* be a binary vector such that $z_i^* = 1, \forall i \in S$ and $z_i^* = 0, \forall i \notin S$. Further let b^* be the optimal solution of the ridge regression in the restricted subspace, i.e.,*

$$
b^* \in \arg\min_{\beta \in \mathbb{R}^p} \left\{ \|X\beta - y\|_2^2 + \rho \|\beta\|_2^2 \;\middle|\; \beta_j = 0, \forall j \notin S \right\}
$$

Then $(b^*, b^* b^{*T}, z^*)$ is optimal to (SDP_{DCL}) if and only if there exists a vector $\tilde{d} \in \mathbb{R}_+^p$ and scalar $\tilde{\lambda} \in \mathbb{R}_+$ such that

$$
\rho^{-1} X^T X + I_p - \mathbf{D}(\tilde{d}) \succeq 0, \tag{7}
$$

$$
\tilde{\lambda} = \tilde{d}_i \left(X_i^T M y\right)^2, \qquad \forall i \in S, \tag{8}
$$

$$
\tilde{\lambda} \tilde{d}_i \geq \left(X_i^T M y\right)^2, \qquad \forall i \notin S \tag{9}
$$

where $M := \left(I_n + \rho^{-1} X_S X_S^T\right)^{-1}$.

The proof of Theorem 2 exploits the optimality conditions of (SDP_{DCL}) and its dual, and is given in detail in the appendix section. We remark that one can directly show that the conditions in Theorem 2 are no stronger than those in Theorem 1.

Remark 1. Suppose that λ is the scalar such that (5) and (6) hold, then (7) – (9) hold for $\tilde{\lambda}$ and $\tilde{d}$, where

$$\tilde{\lambda} := \max\left\{\left(X_j^T My\right)^2 : i \in S\right\}, \tilde{d}_i = \tilde{\lambda}\left(X_j^T My\right)^{-2}, \forall i \in S, \text{ and } \tilde{d}_i = 1, \forall i \notin S.$$

Note that $\tilde{d}_i \in [0,1]$ for all i by construction. Therefore (7) holds. (8) and (9) are also valid by construction.

3 Empirical comparison on exact recovery rate for Gaussian ensemble

In this section we consider the special case of Gaussian ensemble, where the design matrix $X \in \mathbb{R}^{n\times p}$ is generated with i.i.d. N(0,1) entries. A "true" signal β^* is generated to be k-sparse, i.e., it has only k number of nonzero entries, and each nonzero entry is of the order $1/\sqrt{k}$. The response vector y is generated by $y = X\beta^* + \epsilon$, where ϵ has i.i.d $N(0,\gamma^2)$ entries. The following result is established in [2], which characterizes the size of n needed to guarantee the exact recovery of the support of β^* with high probability.

Theorem 3. *There are constants c_0 and c_1, such that the following holds. Suppose that we are given a sample size $n > c_0 \frac{\gamma^2 + \|\beta_S^*\|_2^2}{\beta_{min}^{*2}} \log p$, and that we solve (SDP_{PWG}) with $\rho = \sqrt{n}$. Then with probability at least $1 - 2e^{-c_1 n}$, the relaxation (SDP_{PWG}) is exact, i.e., $v_{PWG} = v_{\ell 0}$.*

Here β_{min}^* is the minimal nonzero entry (in absolute value) of β^*. Note that Proposition 1 ensures that under the same conditions, $v_{DCL} = v_{\ell 0}$ with (at least the same) high probability.

In the remaining part of this section we empirically evaluate the exact recovery for the Gaussian ensemble case. We compare the probabilities of exact recovery by (SDP_{PWG}) and (SDP_{DCL}) for various n and p. To avoid potential numerical issues in solution precision, we exploit Theorem 1 and Theorem 2 to directly search for the certificates. Such a strategy enables us to test whether (SDP_{PWG}) and (SDP_{DCL}) provide the global solution to $(\ell_0 : card)$ on large number of simulated data sets without explicitly solving the semidefinite relaxations many times.

Given simulated data (X, y, β^*), let S denote the support of β^*. Let $\rho > 0$ be fixed, it is straightforward to test whether conditions in Theorem 1 are satisfied. If so, then (SDP_{PWG}) recovers the true support of β^*. The situation is slightly more complicated for (SDP_{DCL}) and Theorem 2. Here we describe a bisection algorithm to search for the dual certificates $\tilde{\lambda}$ and $\tilde{d}$, provided that the support of β^* is used as the index set S.

3.1 A bisection algorithm to search for the dual certificates for (SDP_{DCL})

Without loss of generality we assume that $S = \{1,...,|S|\}$. Firstly if $X_i^T My = 0$ for some $i \in S$, then the convex relaxation (SDP_{DCL}) is not exact unless the trivial case where $X^T My = 0$ for all i. Then without loss of generality we can assume that $\tilde{d}_i = \tilde{\lambda}\left(X_i^T My\right)^{-2}$, for all $i \in S$, and $\tilde{d}_i = \tilde{\lambda}^{-1}\left(X_i^T My\right)^2$, for all $i \notin S$. Therefore the problem of testing (7) – (9) is then equivalent to testing whether there exists $\tilde{\lambda} > 0$ such that the following function is nonpositive,

$$f(\tilde{\lambda}) := \lambda_{\max}\left\{\begin{bmatrix} D_S(\tilde{\lambda}) & 0 \\ 0 & D_{\bar{S}}(\tilde{\lambda}) \end{bmatrix} - \rho^{-1} X^T X - I_p\right\}, \tag{10}$$

where $\lambda\{\cdot\}$ is the largest eigenvalue function, $D_S(\tilde{\lambda})$ is a $|S| \times |S|$ diagonal matrix with diagonal entries $\tilde{\lambda}\left(X_i^T My\right)^{-2}, i = 1,...,|S|$, and similarly $D_{\bar{S}}(\tilde{\lambda})$ is a diagonal matrix with diagonal entries $\tilde{\lambda}^{-1}\left(X_i^T My\right)^2, i = |S|+1,...,p$.

Note that $f(\tilde{\lambda})$ is a convex function when $\tilde{\lambda} > 0$. This is because $\lambda_{\max}\{\cdot\}$ is a convex function and non-decreasing in terms of diagonal entries, and $\tilde{\lambda}^{-1}$ is convex when $\tilde{\lambda} > 0$. It is known that a subgradient of $f(\tilde{\lambda})$ can be computed from an eigenvector associated with the largest eigenvalue in (10). Indeed, let u be such an eigenvector, then

$$h(\tilde{\lambda}) := \sum_{i\in S}\left(X_i^T My\right)^{-2} u_i^2 - \lambda^{-2}\sum_{i\in\bar{S}}\left(X_i^T My\right)^2 u_i^2 \in \partial f(\tilde{\lambda}).$$

In other words, given any $\hat{\lambda} > 0$, $f(\lambda) \geq f(\hat{\lambda}) + h(\hat{\lambda})(\lambda - \hat{\lambda})$ for all $\lambda > 0$. A final ingredient needed for a bisection algorithm is the initial interval. Consider the diagonal entries of the matrix in (10), obviously if $\tilde{\lambda}\left(X_i My\right)^{-2} - (\rho^{-1} X_i^T X_i + 1) \geq 0$ for some $i \in S$, or $\lambda^{-1}\left(X_i My\right)^2 - (\rho^{-1} X_i^T X_i + 1) \geq 0$ for some $i \notin S$, then $f(\tilde{\lambda}) \geq 0$. Therefore we can restrict ourself in a region such that $\tilde{\lambda}\left(X_i My\right)^{-2} - (\rho^{-1} X_i^T X_i + 1) \leq 0$ for all $i \in S$, and $\lambda^{-1}\left(X_i My\right)^2 - (\rho^{-1} X_i^T X_i + 1) \leq 0$ for all $i \notin S$. This provides initial upper and lower bounds such that if there exists $\hat{\lambda}$ such that $f(\hat{\lambda}) < 0$, $\hat{\lambda}$ must be in

$$\left[\max_{i\in\bar{S}}\left\{\left(X_i My\right)^2 (\rho^{-1} X_i^T X_i + 1)^{-1}\right\}, \min_{i\in S}\left\{\left(X_i My\right)^2 (\rho^{-1} X_i^T X_i + 1)\right\}\right].$$

Then the problem of testing whether there is $\tilde{\lambda}$ such that $f(\tilde{\lambda})$ can be solved by the following bisection algorithm,

1. Start with $\ell = \max_{i\in\bar{S}}\left\{\left(X_i My\right)^2 (\rho^{-1} X_i^T X_i + 1)^{-1}\right\}$ and $u = \min_{i\in S}\left\{\left(X_i My\right)^2 (\rho^{-1} X_i^T X_i + 1)\right\}$;
2. Let $\hat{\lambda} = \frac{\ell+u}{2}$ and evaluate $f(\hat{\lambda})$;
3. If $f(\hat{\lambda}) \leq 0$, return YES; otherwise compute $h(\hat{\lambda})$;
4. If $h(\hat{\lambda}) = 0$, return NO. If $h(\hat{\lambda}) > 0$, $u \leftarrow \hat{\lambda} - \frac{\hat{\lambda}}{h(\hat{\lambda})}$; otherwise if $h(\hat{\lambda}) < 0$, $\ell \leftarrow \hat{\lambda} - \frac{\hat{\lambda}}{h(\hat{\lambda})}$. If $u - \ell > \epsilon$, where ϵ is a fixed precision tolerance, then go to step 2. Otherwise return NO.

3.2 Numerical simulations

Using this bisection algorithm, we conduct similar experiments as shown in Figure 1 of [2]. For each value of p (denoted as d in all plots), the true sparsity is set as $\lceil \sqrt{p} \rceil$, and the number of data points $n = \alpha k \log(p-k)$. The true signal $\beta_i^*, (i = 1, \ldots k)$ is generated to be 1 or -1 with same probability. The Figures 1 through 6 show the exact support recovery rate for $\alpha \in [1, 10]$, when ρ are chosen to be $2\sqrt{n}, 3\sqrt{n}, 4\sqrt{n}, 6\sqrt{n}, 8\sqrt{n}, 12\sqrt{n}$.

The numerical simulation illustrates that (SDP_{DCL}) can recover the support of true signals with significant less data points than that of (SDP_{PWG}). Also the exact recovery rate of (SDP_{DCL}) appears to be much less sensitive to the choice of ρ. This result further motivates us to study scalable approximate methods, such as those based on low rank factorization of the matrix B, to solve (SDP_{DCL}).

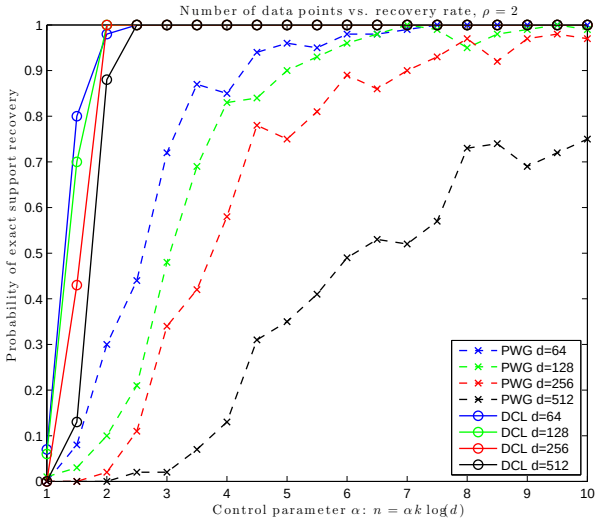

Figure 1: Exact support recovery rate when $\rho = 2\sqrt{n}$.

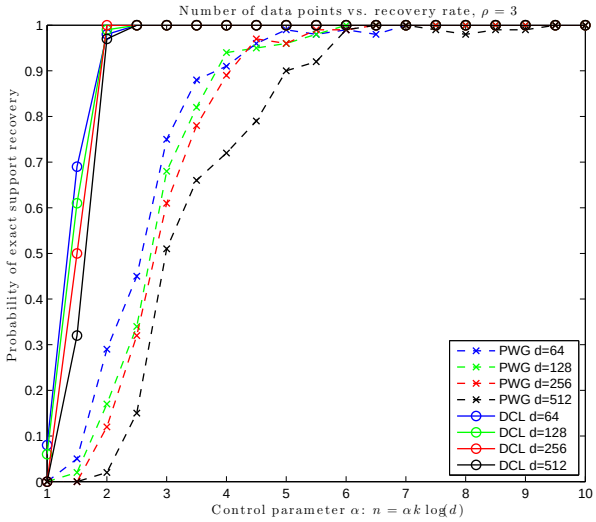

Figure 2: Exact support recovery rate when $\rho = 3\sqrt{n}$

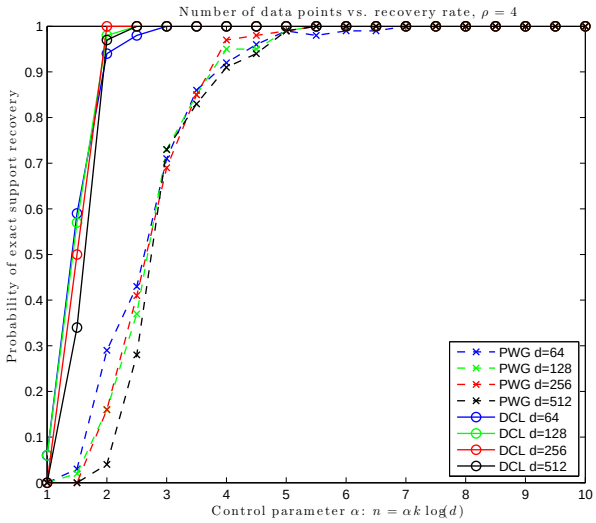

Figure 3: Exact support recovery rate when $\rho = 4\sqrt{n}$

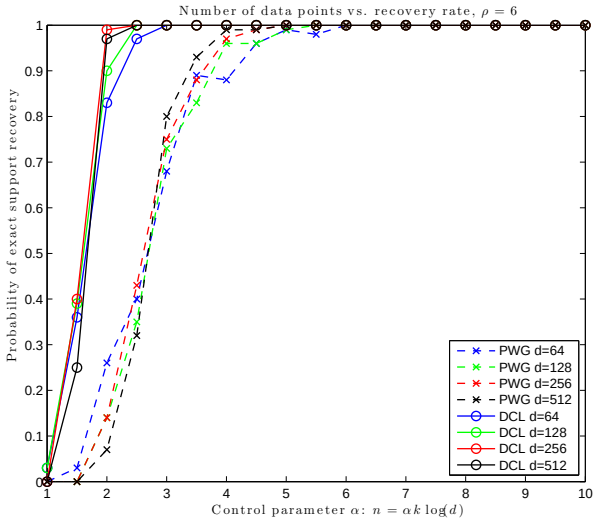

Figure 4: Exact support recovery rate when $\rho = 6\sqrt{n}$

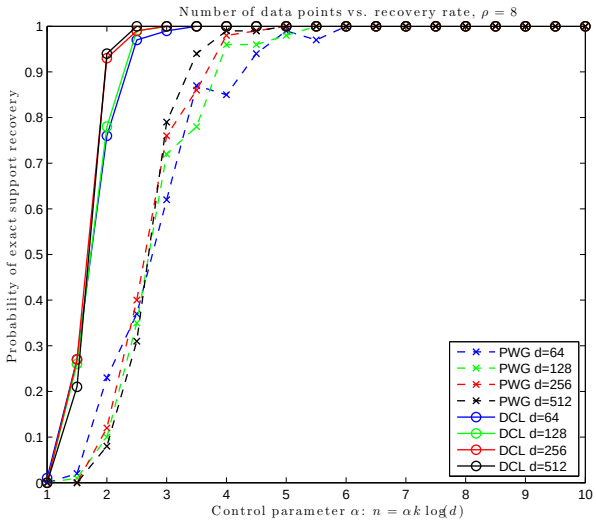

Figure 5: Exact support recovery rate when $\rho = 8\sqrt{n}$

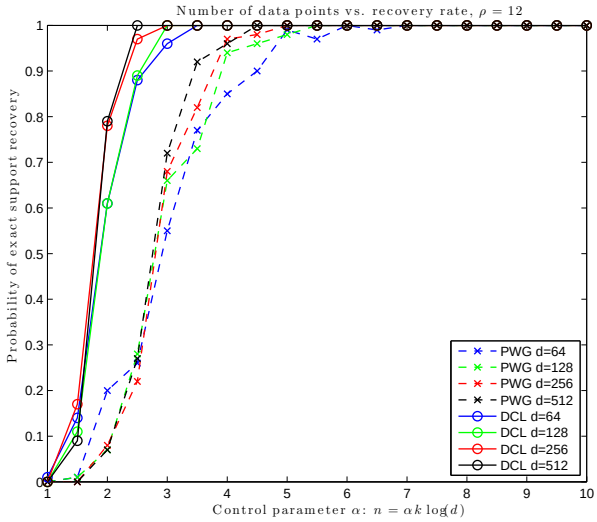

Figure 6: Exact support recovery rate when $\rho = 12\sqrt{n}$

References

[1] Hongbo Dong, Kun Chen, and Jeff Linderoth. Regularization vs. Relaxation: A conic optimization perspective of statistical variable selection. Submitted to Math. Prog. A, 2015.

[2] Mert Pilanci, Martin J. Wainwright, and Laurent El Ghaoui. Sparse learning via Boolean relaxations. *Mathematical Programming (Series B)*, 151:63–87, 2015.

Appendix

Proof of Theorem 2.

Proof. Without loss of generality we assume $S = \{1, ..., k\}$. When $z = z^*$, the constraints in (SDP_{DCL}) enforce that $B_{ii} = b_i = 0, \forall\, i \notin S$. Further since z^* is binary, we may assume $B^* = b^* b^{*T}$ without loss of generality, where b^* solves the restricted regression problem:

$$b_S^* \in \arg\min_{\beta \in \mathbb{R}^{|S|}} \left\| X_S \beta - y \right\|_2^2 + \rho \|\beta\|^2 \;\; \Rightarrow \;\; b^* = \begin{bmatrix} \left(\rho I + X_S^T X_S\right)^{-1} X_S^T y \\ \mathbf{0}_{|S^c|} \end{bmatrix}$$

and $\mathbf{0}_{|S|}$ is a $|S| \times 1$ zero vector. By strong duality and the KKT conditions, there exists b^*, B^* such that (b^*, B^*, z^*) is optimal in (SDP_{DCL}) if and only if there exists dual variables (τ, λ, t, d) such

that the first order optimality condition holds

$$\begin{bmatrix} \tau & -y^T X - t^T \\ -X^T y - t & X^T X + \rho I - \mathbf{D}(d) \end{bmatrix} \succeq 0 \tag{11}$$

$$\begin{bmatrix} \lambda & t_i \\ t_i & d_i \end{bmatrix} \succeq 0, \forall i \tag{12}$$

$$\begin{bmatrix} \tau & -y^T X - t^T \\ -X^T y - t & X^T X + \rho I - \mathbf{D}(d) \end{bmatrix} \bullet \begin{bmatrix} 1 & b^{*T} \\ b^* & b^* b^{*T} \end{bmatrix} = 0 \tag{13}$$

$$\begin{bmatrix} \lambda & t_i \\ t_i & d_i \end{bmatrix} \bullet \begin{bmatrix} 1 & b_i^* \\ b_i^* & (b_i^*)^2 \end{bmatrix} = 0, \ \ \forall i \in S \tag{14}$$

In (14), only complementarity conditions for $i \in S$ are needed because for all $j \notin S$, z_j^* equals 0, which implies that $b_j^* = 0$ by feasibility and thus the complementarity condition holds. Now we aim to derive simpler conditions on the existence of such dual variables (τ, λ, t, d). With condition (11), (13) can be equivalently written as

$$\begin{bmatrix} \tau & -y^T X - t^T \\ -X^T y - t & X^T X + \rho I - \mathbf{D}(d) \end{bmatrix} \begin{bmatrix} 1 \\ b^* \end{bmatrix} = 0,$$

which is further equivalent to the following two equations (15) and (16),

$$t = (X^T X + \rho I - \mathbf{D}(d)) b^* - X^T y \tag{15}$$

$$\tau = b^{*T} t + y^T X b^* = b^{*T} (X^T X + \rho I - \mathbf{D}(d)) b^*. \tag{16}$$

Now we exploit (15) and (16) to eliminate t and τ in (11),

$$\begin{bmatrix} \tau & -y^T X - t^T \\ -X^T y - t & X^T X + \rho I_p - \mathbf{D}(d) \end{bmatrix} = \begin{bmatrix} b^{*T} \\ I_p \end{bmatrix} \left[X^T X + \rho I - \mathbf{D}(d) \right] \begin{bmatrix} b^* & I_p \end{bmatrix}.$$

Therefore conditions (11) and (13) are equivalent to (15), (16) and

$$X^T X + \rho I_p - \mathbf{D}(d) \succeq 0. \tag{17}$$

Now we consider conditions (12) and (14). Again with (12), (14) is equivalent to

$$\begin{bmatrix} \lambda & t_i \\ t_i & d_i \end{bmatrix} \begin{bmatrix} 1 \\ b_i^* \end{bmatrix} = 0, \forall i \in S \iff t_i = -d_i b_i^*, \lambda = -t_i b_i^* = d_i (b_i^*)^2, \forall i \in S$$

We claim that for all $i \in S$, $t_i = -d_i b_i^*$ is implied by (15). Indeed, as b^* minimizes the convex quadratic form in the restricted subspace corresponding to S,

$$\begin{aligned} 0 = \frac{d}{d\beta_i}\bigg|_{\beta = b^*} \|X_S \beta - y\|_2^2 + \rho \|\beta\|^2 &= 2 X_i^T (X_S b^* - y) + 2\rho b_i^*, \quad \forall i \in S, \\ &= 2 \left(X_i^T X b^* - X_i^T y + \rho b_i^* \right). \end{aligned}$$

So the i-th row of (15) can be equivalently written as,

$$t_i = X_i^T(Xb^* - y) + \rho b_i^* - d_i b_i^* = -d_i b_i^*, \ \ \forall i \in S.$$

Therefore conditions (12) and (14) can be simplified as,

$$\begin{aligned} d_i &\geq 0, \qquad \forall i = 1, ..., p, \\ \lambda &= d_i (b_i^*)^2, \qquad \forall i \in S, \\ \lambda d_i &\geq t_i^2, \qquad \forall i \notin S. \end{aligned}$$

Note that for all $i \notin S$, $b_i^* = 0$. So by (15), $t_i = X_i^T(Xb^* - y)$ for all $i \notin S$. Therefore the optimality conditions (11) – (14) are equivalent to

$$(15),\ (16),\ (17),\ \lambda = d_i (b_i^*)^2, \forall i \in S,\ \lambda d_i \geq \left(X_i^T(Xb^* - y)\right)^2, \forall i \notin S.$$

Note that (15) and (16) simply state that and t and τ are uniquely determined once d is fixed, where t and τ do not appear in other conditions. To complete the proof it suffices to prove two sets of equalities:

$$b_i^* = \rho^{-1} X_i^T (Xb^* - y), \qquad \forall i \in S, \tag{18}$$

and

$$-X_i^T(Xb^* - y) = X_i^T(\rho I + X_S X_S)^{-1} y, \qquad \forall i. \tag{19}$$

Our conclusion then follows after a rescaling $\tilde{d} = d_i \rho^{-1}$ and $\tilde{\lambda} = \lambda \rho^3$. Indeed, the equalities (18) and (19) can be proved by using the Sherman-Morrison-Woodbury formula,

$$\begin{aligned} X_i^T(Xb^* - y) &= X_i^T \left(X_S \left(\rho I + X_S^T X_S \right)^{-1} X_S^T y - y \right) \\ &= -X_i^T \left(I - X_S \left(\rho I + X_S^T X_S \right)^{-1} X_S^T \right) y \\ &= -X_i^T \left(I + \rho^{-1} X_S X_S^T \right)^{-1} y; \qquad \forall i \end{aligned}$$

$$\begin{aligned} b_S^* &= \left(\rho I + X_S^T X_S \right)^{-1} X_S^T y \\ &= \left[\rho^{-1} I - \rho^{-2} X_S^T \left(I + \rho^{-1} X_S X_S^T \right)^{-1} X_S \right] X_S^T y \\ &= \rho^{-1} X_S^T \left(I + \rho^{-1} X_S X_S^T \right)^{-1} y. \end{aligned}$$

□